

Advice to Authors of Extended Abstracts

William Pugh

Dept. of Computer Science and Institute for Advanced Computer Studies

Univ. of Maryland, College Park

This article stems from discussions among the program committee for SIGPLAN'91 PLDI. The program committee thought it might be useful to put together some advice for authors. To give some context to these suggestions, I've also provided a brief description of the process by which the conference papers were reviewed, partially from my perspective. This process is similar to the way most SIGPLAN conferences are run, although the details differ for each conference.

How the Papers Were Reviewed:

There were 169 extended abstracts submitted to the SIGPLAN '91 PLDI conference. At the request of the program committee chair, program committee members (and their graduate students) refrained from submitting any abstracts to the conference. This allowed us to avoid having to deal with direct conflict of interests.

Each program committee member was assigned 60 abstracts, based on his or her areas of expertise. Since all abstracts were sent to all committee members, members could review any abstracts they wished, so long as they reviewed at least the abstracts assigned to them. Program committee members could review the abstracts themselves or have others review them, although in most cases the program committee members at least briefly reviewed all the abstracts they were assigned, even if they had colleagues review some of them in detail for them.

I wasn't able to read any abstracts until after the semester ended in mid-December, and I allowed myself a week off from reading abstracts for Christmas. Thus, I had about four weeks to read the abstracts, and I couldn't spend much more than 20 hours a week reading them (due to limitations both on available time and the amount of reviewing I could do in a day before I suffered burnout). Since I read more abstracts than I was assigned, this came down to an average of one hour per abstract.

In reading an abstract, I had to try to understand the work presented, the significance of it, and possible problems with it. I spent at least 30-40 minutes on almost every abstract, sometimes coming back to an abstract several times. I spent over four hours each on several abstracts. In one case this was because the abstract looked interesting but was badly written; in another case, because the abstract dealt with a dense subject. In several cases, I spent several hours on a paper simply because I had expertise or interest in the topic described by the paper.

The program committee met for two days to discuss the submitted abstracts and choose those to be accepted. A preliminary numerical ranking provided by the reviews received in advance of the meeting helped structure our discussions. On each of several passes through all the submissions, some papers were eliminated from consideration, others were

retained for further discussion and some were accepted. Finally, we had a total of 28 accepted papers.

What is an Extended Abstract?

An extended abstract is not simply a long abstract. An extended abstract should contain references, comparisons to related work, proofs of key theorems and other details expected in a research paper but not in an abstract.

An extended abstract is a research paper whose ideas and significance can be understood in less than an hour. Writing an extended abstract can be *more* demanding than writing a research paper.

Some things that can be omitted from an extended abstract: future work, details of proofs or implementation that should seem plausible to reviewers, ramifications not relevant to the key ideas of the abstract.

Some Issues Considered by the Committee:

- **Are there any major technical flaws in the abstract?** In a few rare cases, reviewers found serious technical flaws in a submission.
- **Is the work a significant advance over previous work in the area, by the same authors or others?** The abstract should give a clear description of the advantages offered by the new technique over previous techniques. Simply describing an interesting new way of doing something that could be done as simply and efficiently by previous techniques won't get an abstract accepted. The best abstracts gave a clear description of what their results allowed that couldn't be done previously and why that is significant. Examples and measurements are great for this.

A related problem is not citing relevant work in the area. Don't rely on the program committee realizing that X's work in this area doesn't apply because you are considering a slightly different problem that renders X's methods unusable.

If you have additional current papers on topics related to your submission (accepted by or submitted to other conferences or journals), be sure to discuss the contribution of your submission over that of your other papers.

- **If the work involves a specialized application, does it make a more general contribution?** Some abstracts described interesting specialized applications. Much of the content of these abstracts involved descriptions of the context of the work or applying standard techniques in the new context. In some cases, it was unclear if the resulting paper would be useful to people not interested in the author's specific application.

If you submit an extended abstract involving a specialized application, be sure the significant contributions of your work don't get lost in the details of your application.

- **Does the abstract offer an interesting perspective on a problem or describe experience that might be useful to others?** Several committee members lamented that although several authors had built substantial systems, and tried several approaches to learn which ones worked and which ones didn't, the authors only

wrote abstracts about narrow technical results related to their systems. Relevant comments about practical experiences attempting to apply new technologies can significantly increase the value of any paper.

- **Is the abstract well presented and understandable?** We didn't reject any abstracts for being poorly presented. However, all other things being equal, the program committee was more enthusiastic about abstracts that were clear and well presented.
- **Is the abstract too long?** There are many methods of trying to fit 20 pages of material into the 10 page limit on extended abstracts (reducing margins, using 9-point type on 10-point leading with double columns, etc.) **They are all strongly discouraged.** The page limit is to encourage authors to write abstracts that can be absorbed quickly, not to save trees, (although our request for double-sided copies of the abstracts did have this as its goal). No abstracts were rejected purely for reasons of length, but none of the accepted abstracts significantly violated the spirit of the 10 page limit; consider this a strong hint.

Several program committee members stated that after reading 10 pages worth of material, they felt free to stop reading at any point if they were not truly excited by the paper.

Don't let the page count limit prevent you from providing figures or examples that make the paper easier to understand. The page count limit should be considered an upper bound on the number of full pages of text, exclusive of figures and examples. One program committee member disagreed and thought that the page limit should be strictly adhered to, noting that if a picture is worth 1000 words, a picture is worth more than the 200 words it displaces.

In exceptional cases, it may be appropriate to put additional material in an appendix that extends past the length limit. This is acceptable *only if* the extended abstract itself stands on its own without the additional material. Given their time limitations, most reviewers probably *will ignore* the appendix. Acceptable material for an appendix could include background material for committee members not familiar with the details of the research area and details of proofs and implementations omitted from the body of the abstract.

- **Does the abstract address the obvious questions raised by the research?** For example, if an abstract claims to describe "an efficient, practical algorithm" for something, it should give empirical timings, asymptotic analysis or both. If the techniques described require solving a problem that is NP-Complete or undecidable in general, the abstract should discuss the difficulty of solving the problem. It may be that in practice the problems that arise in the author's application can be solved efficiently; but if the abstract doesn't discuss it, the program committee doesn't know if the author is even aware of the potential problem.

The program committee was sympathetic about not expecting data that ought to have been very difficult to collect. However, the committee was disappointed in several instances by abstracts that failed to report data that ought to have been easy to collect and would have answered obvious questions about the work.

Final Comments for Authors:

- An ideal submission should have a reviewer intrigued within the first 5 minutes of reading, excited within 15 minutes and satisfied within 45 minutes. If your abstract fails any of these tests, it might be rejected no matter how good the research is. Committee members may spend more than 30-45 minutes on your abstract, but you shouldn't rely on it.

Before you submit an abstract, give it to a programming languages colleague who is not familiar with the details of your research or your research area and ask him or her how much they can get out of it in less than an hour.

- Don't overlook the importance of the introduction, figures, examples, and conclusions (and measurements if applicable) in an extended abstract.
- Remember that some program committee members, of necessity, are not experts in your area of research and that when they pick up your abstract they may have already reviewed 8 abstracts that day. Material that may take an expert in your area 5 minutes to go through might take some committee members 20 minutes or more.
- There are some types of research that are difficult to publish in a conference simply due to the amount of time and effort that would be required for the program committee members to review the abstract properly. If you can't prepare an extended abstract of your work that can be digested and its significance understood in an hour, it may not be possible to accept your paper, no matter how good the described research. For some types of research (particularly research on new topics), it may be impossible to meet this standard, no matter how well you write. This is an unfortunate flaw in the system, and we have no remedy except to suggest that you submit your paper to a journal where more time can be taken to referee it properly.
- Please remember that we cannot give as much attention to a submission as would be given to a journal submission, and we do make mistakes. If you get back comments that suggest the program committee misunderstood your abstract, you can use that information constructively. If the program committee misunderstood your work, other readers may misunderstand it as well.
- This note has placed a lot of emphasis on the idea that an extended abstract need to be clearly written and easy to understand. Of course, whenever possible that standard should be applied to full papers as well.
- For additional advice, read the excellent article by Mark Wegman that inspired this report: "What it's like to be a POPL referee; or How to write an extended abstract so that it is more likely to be accepted," SIGPLAN Notices, Vol. 21, No. 5, May 1986, pages 91-95.